RESYNTHESIS ON THE SYNCLAVIER (R)

THE SFM ANALYSIS PROGRAM

A New Software Module for
the Sample-to-Disk (TM) System

July 1984

Preliminary Version

TABLE OF CONTENTS

PREFACE

The New England Digital Resynthesis System provides the tools for a powerful, new synthesis approach. Heretofore, the synthesis of acoustic instrument sounds has been considered to be beyond the capabilities of any real-time synthesizer. Now, with the Sample-to-Disk (TM) Analysis Program used along with the Resynthesis Release of the Synclavier (R) Real-Time Performance System (Release R), you can synthesize life-like acoustic instrument timbres as well as a range of new timbres unrealizable on any other digital synthesizer. Intuitive in its approach, the new system reduces timbre development time and expands your synthesis horizons to new frontiers.

The process of Analysis and Resynthesis can be likened to the process of filming a scene with a movie camera and then playing back the scene through a movie projector. The movie camera captures a series of static images, or frames, which are pictures of the scene at distinct moments in time. When replayed through the movie projector, the different frames combine to produce the illusion of smooth movement and convey a realistic representation of the original scene. In a similar fashion, the Analysis Program captures the acoustic properties of a sound at a number of discrete points in time. These static "pictures" of the sound are then recombined by the Resynthesis Real-Time Performance system so as to portray convincingly much of the character of the original sound.

The Analysis Program enables you to accurately measure the pitch, spectrum, and volume of a sampled sound at various points in time. Unlike methods based upon the Fast Fourier Transform (FFT), which are unable to accurately measure the frequency components in fast changing waveforms because they require long samples to achieve high resolution in the frequency domain, this method is able to accurately measure the harmonic content of waveforms on a cycle by cycle basis.

Once these measurements have been made, the Analysis Program converts them into a timbre definition with a series of timbre frames. Then, Resynthesis Release of the Synclavier (R) Real-Time Performance system is accessed, the new timbre is activated, and the sound may be played from the keyboard or guitar, as well as modified and recorded just like any timbre.

This manual describes the Analysis Program. The partner manual Resynthesis on the Synclavier (R) - Timbre Frame Sound Construction describes the Synclavier (R) Real-Time part of the system. You should read that manual first, so that you understand the basics of timbre frame programming.

PREPARATION OF SOUND FILES FOR ANALYSIS

Before using the Analysis Program, you must prepare a sound file by adding pointers, or labels, to each part of the file you wish to emulate in a timbre frame. The idea is to create just enough timbre frames to capture the essence of the sound and no more. If too many frames are used, the efficiency of the Synclavier (R) Real-Time system will be reduced. If too few frames are used, the resynthesized sound will not be faithful to the original sampled sound.

The range of sounds that you might want to resynthesize is not restricted to a few well-understood instrument sounds. Therefore, it would be both difficult and inefficient for the computer to decide where each new "photograph" should be taken. Since you will manually place labels at the locations in the sound which you deem important to the identity of the sound, you can maintain precise control over the analysis. In this way, your own expertise as a synthesist can be readily incorporated into the resynthesis process.

Sound file preparation is accomplished through the Signal File Manager (SFM). You call up the sound file you wish to resynthesize, using the OLD command. Then you move the cursor to the locations you think are relevant to the overall perception of the sound. Using the LABEL command, you then place labels at the beginning and end of each cycle of the waveform which seems to be "important" to the sound. (A more powerful technique will be described below which will activate the automatic cycle extractor; this technique requires the placement of only one label per frame.)

To maintain a distinction between labels used as pointers for the Analysis Program and those used for other purposes, a special format for Analysis labels has been adopted. Analysis labels must be suffixed with the characters "_1" or "_2". A label suffixed with "_1" denotes the start of a fundamental cycle; a label consisting of the same letters suffixed with "_2" denotes the end of that cycle. During the Analysis, the samples between the two labels (inclusive of endpoints) will be decomposed into their constituent harmonics. This information will create the waveform for the related frame in the Synclavier (R) timbre definition. The time between one "_1" label and the next "_1" label will be used to set the cross-fade time for the timbre definition.

Some examples of Analysis labels, as well as ordinary labels, are given below:

Analysis Labels:

A_1	APPLE_1	SND_1	START_1	LABEL_1
A_2	APPLE_2	SND_2	START_2	LABEL_2

Ordinary Labels:

ORIGIN	END	ATTACK	P1	P2
A	SND	TROMBONE	FLUTE	AXE

Examining A Sound File

Call up the demonstration sound file HORNF4 stored on the Analysis Demonstration Timbre diskette. It has been labeled as an example for study. You can see that one cycle occurring during the attack of the trumpet sound has been specified by labels "A_1" and "A_2." Note that label "A_2" has been placed one sample to the left of the zero-crossing of the cycle since the sample at the zero-crossing represents the first point in the next cycle.

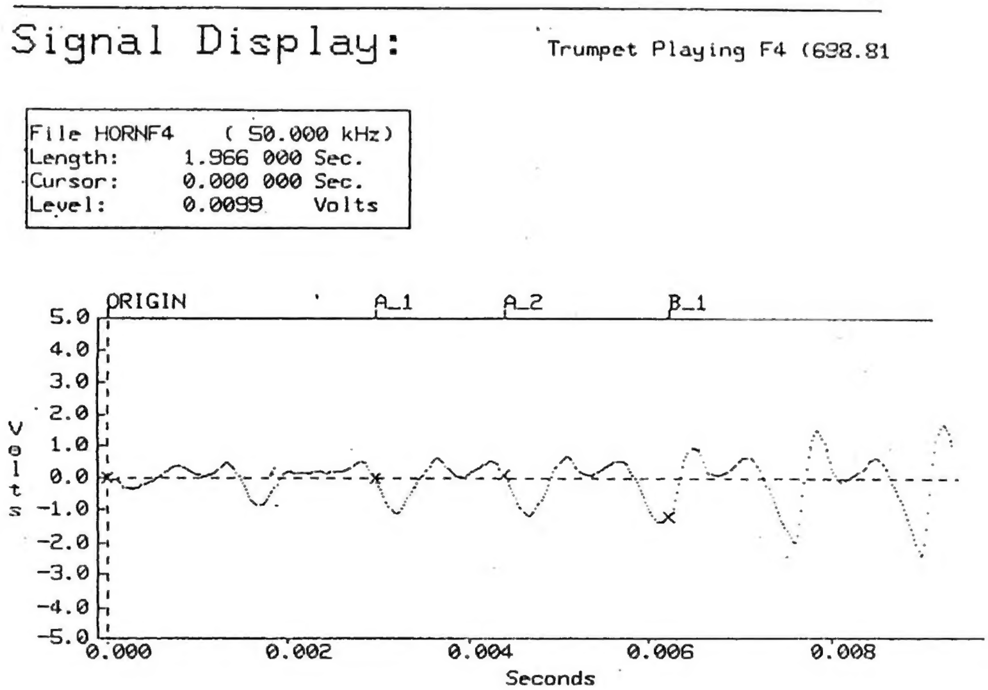

In order to minimize the number of labels which need to be placed in a sound file, a feature has been implemented in the Analysis Program which can automatically extract complete cycles given only the desired starting point for the cycle. This feature performs very well on sounds having a relatively constant pitch and a relatively harmonic spectrum. To activate the feature, you need label only the starting point of the desired cycle. When the Analysis Program sees a label with the suffix "_1" but finds no matching label suffixed with "_2", it will automatically determine the best cycle length to use.

Since the HORNF4 sound file has a relatively constant pitch and a relatively harmonic spectrum, it is an excellent candidate for automatic cycle extraction. The label "B_1", since it is not followed by "B_2", will activate the automatic cycle extractor. You can now use the UNLABEL command to remove the "A_2" label from demonstration sound file HORNF4, so that the automatic cycle extractor will be invoked for the first cycle as well.

During the Analysis, the samples in the cycle starting with A_1 will be analyzed and converted into a set of harmonic coefficients that will be used to compute the waveform for Frame 0; the time between the origin and A_1 will be used to set the VOLUME ENVELOPE ATTACK.

The samples in the cycle starting with B_1 will be converted into the coefficients for Frame 1 and the time between A_1 and B_1 will be used to set the splice time for Frame 1.

Spend some time looking at the positioning of the other labels in HORNF4. They have been spaced very close together during the attack portion of the sound. As the sound evolves, they are placed at progressively greater intervals. In most cases, you will find that this sort of labeling scheme will result in the most desirable timbral transitions.

As the example trumpet note attacks, the upper harmonics evolve at a very quick rate. Many frames are needed during this period in order to capture all the changes. As the note progresses, the spectrum changes at a increasingly slower rate. During this period, the number of frames used can be gradually reduced.

Note that the Synclavier (R) Real-Time Performance system performs most functions on a five-millisecond time base. In theory, it is possible for you to position labels so that adjacent frames are less than five milliseconds apart. The Analysis Program would have no problem analyzing such a file. However, the Synclavier (R) Real-Time Performance system will not be able to step through frames any faster than one frame per five-millisecond time period. If multiple frames have crossfade times of less than five milliseconds, the Synclavier (R) Real-Time system will play back the frames at a slower rate than specified in the timbre definition. In all probability, you would not notice this slowdown.

ACTIVATING THE ANALYSIS PROGRAM

Once a sound file has been labeled, you are ready to begin the analysis procedure. The Analysis Program can only be invoked from within the Signal File Manager, using the SFM command, SYN(nthesize). The sound file that is current in the SFM will be the file to be analyzed. Therefore, with your labeled HORNF4 as the current sound file, type SYN.

After a brief pause, the screen will clear, and the New England Digital Analysis Program Main Menu will be displayed. In the uppermost box, you will see the name of the program, the release date of your version, the maximum number of timbre frames the Analysis Program can support, and the name of your current timbre. The left-hand box will contain seven system commands and the right-hand box will contain six analysis parameters.

Before we describe the system commands and the analysis parameters, you should be aware that a sound has three different forms in the Analysis Program. Here is a quick overview of how the different forms interact and how you save and recall each one.

The first form is the original labeled sound file. You recall and save these in the SFM using the OLD and SAV commands. The second form is the Analyzed timbre. You convert the sound file into an Analyzed timbre by pressing KP1 from the Analysis Program Main Menu. You can save and recall Analyzed timbres from the Analysis Main Menu using the KP2 and KP3 commands. The third form is the Synclavier (R) timbre. You convert the Analyzed timbre into a Synclavier (R) timbre by pressing PF2. You can perform, save, and recall Synclavier (R) timbres from the keyboard.

USING THE SYSTEM COMMANDS

You will use the keys on the keypad (KP) as well as the BREAK key to issue commands:

<KP1> - Begin Analysis of the Current Sound File

Press KP1 to begin the analysis of the current SFM sound file. If there is already an Analyzed timbre in the frame memory from an earlier analysis or from an Analyzed timbre recall, you will be asked to verify this command by pressing RETURN before the frame memory will be cleared and a new analysis will be initiated.

<KP2> - Save an Analyzed Timbre

Press KP2 to save an Analyzed timbre. You will then be asked to type in a filename for the Analyzed timbre. Unless there is no room on your Winchester disk, the data will then be stored in the specified file. The save timbre command will store the following information:

1. The name of the sound file from which the data came.
2. The number of frames in the Analyzed timbre.
3. The number of harmonic coefficients per frame.
4. The harmonic coefficients and phases for each frame.
5. The pitch of each frame.
6. The volume of each frame.
7. And a record of all errors which occurred during analysis of the sound file.

<KP3> - Recall an Existing Timbre

Once an Analyzed timbre has been saved using the above command, it may be recalled by pressing KP3. You will then be asked to provide the filename of a valid Analyzed timbre. If the specified filename is found in your current catalog, the timbre will be read into memory, replacing the current timbre, if there is one. The name of the sound file from which the recalled timbre was created will be written at the top of the screen.

<KP4> - Display Errors from Analysis of Current Timbre

The KP4 command will display any errors which occurred during the analysis of the current Analyzed timbre. There are currently three kinds of analysis errors which might occur.

A PITCH ERROR occurs whenever the automatic cycle extractor is unable to locate a satisfactory cycle with the given APPROXIMATE PITCH OF SOUND and PITCH WINDOW parameters (see below). Perhaps you were off by an octave in your pitch specification, or the sound changed during the sound file more than you thought. You can go back to the Analysis Main Menu and change either of these parameters.

A BANDWIDTH ERROR occurs when the NUMBER OF HARMONICS parameter is greater than one-half the number of samples in the extracted cycle period. For example, in order to extract 24 harmonics from a cycle, the cycle must be at least 48 samples long. If the extracted cycle has a period of less than 48 samples, a BANDWIDTH ERROR will result and the number of harmonics extracted will be limited to one-half of the extracted cycle length. This is really a warning, rather than an error.

A LENGTH ERROR will occur if a cycle to be analyzed exceeds 2048 samples in length. At present, the Analysis Program can not process more than 2048 samples per frame. If this error occurs, the frame will not be analyzed. The program will proceed to the next frame. This will probably never occur if you use automatic cycle extraction, but only if the segment of the file between a pair of labels contains too many samples.

<PF2> - Set Up Timbre and Overlay to Synclavier (R) Real-Time System

Press PF2 to transfer control to the Synclavier (R) Real-Time system. After you enter this command, the program will pause for a time while the Analyzed timbre is converted into the Synclavier (R) timbre format. Upon completion of the timbre conversion process, the Synclavier (R) Real-Time Performance system will be activated. The new Synclavier (R) timbre will be active on the keyboard and may be saved on the disk or modified further.

If you press PF2 before performing any analysis or before recalling an existing Analyzed timbre, the Analysis Program will transfer control to the Synclavier (R) Real-Time Performance system as usual but will set up a simple sine wave timbre on the keyboard.

<PF4> - Return to SFM

Press PF4 to quit the Analysis Program and return to SFM.

<BREAK> - Return to the Monitor

Press BREAK to quit the Analysis Program and return to the Monitor.

SETTING THE ANALYSIS PARAMETERS

There are six analysis parameters which modify the way in which the analyzed data is interpreted by the program. Four of these parameters are set by moving the cursor to the line of the desired parameter, typing in the desired value, and then pressing RETURN. The other two parameters either enable or disable special functions of the Analysis Program. You move the cursor to either of these parameters and press any key on the keyboard to toggle between the two modes.

APPROXIMATE PITCH OF SOUND

This parameter is used to control two aspects of an analysis: automatic cycle extraction and calculation of pitch envelopes.

When the automatic cycle extractor is activated (by labeling only the starting point of the desired cycle), it is necessary for the system to know the approximate pitch of the sound in order to properly locate one complete cycle of the fundamental. The automatic cycle extractor uses the APPROXIMATE PITCH OF SOUND value, in conjunction with the PITCH WINDOW parameter, described below, to aid in determining the exact period of the wave.

When the USE RAW PITCH mode has been enabled, the APPROXIMATE PITCH OF SOUND parameter is also used to as a reference when determining the appropriate PARTIAL TUNING value for the resynthesized Synclavier (R) timbre.

The APPROXIMATE PITCH OF SOUND parameter is expressed in terms of its octave.pitchclass number, like the PITCH and OCTAVE parameters on the SFM SET menu. Allowed values for the APPROXIMATE PITCH OF SOUND parameter range from a low of 0.0000 (about 32.7 Hz) to a high of 9.0307 (about 20 kHz).

The default value for APPROXIMATE PITCH OF SOUND is obtained from the OCTAVE value stored with the sound file.

PITCH WINDOW

The PITCH WINDOW parameter is used by the automatic cycle extractor in conjunction with the APPROXIMATE PITCH OF SOUND parameter. When you want the Analysis Program to determine the endpoints of wave cycles automatically, it is necessary to know both the approximate pitch of the sound as well as the range of cycle periods over which the computer should search. The PITCH WINDOW controls the range of cycle periods that should be examined. This parameter is expressed in terms of percent variation from the APPROXIMATE PITCH OF SOUND parameter and may run from zero percent (no deviation) up to 50 percent deviation.

The default PITCH WINDOW setting is 20 percent.

> Example: You have a recording of a trumpet made with a 50 kHz sampling rate. If the approximate pitch of the sound was about 440 Hz, you would set APPROXIMATE PITCH OF SOUND to 3.0900 (A3). At a 50 kHz sampling rate, one cycle of the trumpet wave would be approximately 114 samples long (50 kHz/440 Hz = 113.636 samples). By setting PITCH WINDOW to 10 percent, you would force the cycle extractor to look at cycles having periods of from 103 samples to 125 samples (114 samples +/- 10%) in choosing the best cycle length for the labeled wave.

If the cycle extractor should be unable to find a satisfactory cycle length with a particular combination of PITCH and WINDOW settings, the extractor will use a cycle period corresponding to the APPROXIMATE PITCH OF SOUND setting. In general, this will never happen unless you have used too narrow a PITCH WINDOW value. You can see that if your sound were to vary widely in its pitch, you would want to choose a larger pitch window. Conversely, if the pitch were very constant, you would be able to choose a smaller window.

Since it takes longer for the computer to extract cycles when the PITCH WINDOW is large, it is desirable to choose a window which is as narrow as possible. However, if the PITCH WINDOW is set too small, the cycle extractor may not be able to accurately extract a cycle of the fundamental. This condition would result in an error and would mean that the frame on which the error occurred was not necessarily an accurate representation of the original waveform.

USE RAW PITCH

The Analysis Program is able to establish a pitch envelope for the sound by computing a pitch offset for each frame, if desired. This pitch envelope captures any changes of pitch which occurred in the original sound. By setting the USE RAW PITCH parameter to "YES", you will enable this feature. The cycle length of each frame will be used to compute a pitch glide (or offset), in semitones, that should occur as one frame fades to the next. Since Frame 0 can have no pitch offset value, the cycle length for Frame 0 is used to compute a PARTIAL TUNING value. This PARTIAL TUNING will differ from 440.0 Hertz by as many semitones as the pitch of Frame 0 differed from the APPROXIMATE PITCH OF SOUND parameter. If USE RAW PITCH is set to "NO", no pitch envelope will be computed, and the PARTIAL TUNING will be set to 440.0 Hertz.

The default setting for USE RAW PITCH is "NO" (disabled).

NUMBER OF HARMONICS

The NUMBER OF HARMONICS parameter determines how many harmonic coefficients will be extracted from each frame in the sound file. In previous versions of the Synclavier (R) Real-Time system, one was limited to 24 harmonic coefficients per wave. Using this new analysis method, you may extract, and use for your Synclavier (R) timbre, up to 128 coefficients for each frame!

Before you start the Analysis Program, you will want to look at the spectrum of the sound you are going to resynthesize using the SPECTRUM command in SFM. The more harmonics you extract from each frame, the longer the analysis will take. If a sound only has 32 harmonics, it would be a waste of time to extract all 128 harmonics, since the upper 96 harmonics would always have zero amplitude.

The default setting for NUMBER OF HARMONICS is 24.

> Special Note: If you choose to analyze for more than 24 harmonic coefficients, you will not be able to modify those harmonic coefficients from the Synclavier (R) Real-Time Performance program. The reason for this temporary limitation is that, presently, only the first 24 harmonic coefficients are stored in the computer as part of the timbre definition. If the waveforms should need to be recomputed because you have dialed in a new value for one of the coefficients, frequency components above the 24th harmonic will be lost. This limitation will be eliminated in a future release of the Synclavier (R) Real-Time system.

USE RAW PHASES

A new feature in the Resynthesis Release of the Synclavier (R) Real-Time system is the ability to control the phase relationships between individual harmonics in a waveform. By setting the USE RAW PHASES parameter to "YES", the Analysis Program will preserve the phase relationships extracted from the waveforms of each frame. Enabling this feature will result in resynthesized waveforms which are visually very similar to the waveforms in the original sound file. If USE PHASE INFORMATION is set to "NO", the phase relationships between harmonics will be ignored. All harmonics will be reconstructed using zero phase.

A few points regarding the use of phase in the Analysis Program should be clarified:

1. When listening to individual waveforms, the phase relationships between the harmonics will not affect the tone of the sound. If you dial in a set of harmonics and then change the phases of some of the harmonics, the waveform will still sound the same.

2. If the USE RAW PHASES mode is not enabled, you will be able to locate labels in the sound file without regard to zero crossings. Since all phase information will be disregarded in the reconstruction of the original waveforms, the phase of your labels relative to the phase of the waveforms will be unimportant. This mode simplifies the labeling process significantly.

3. When the USE RAW PHASES mode is disabled, the Analysis Program optimizes the analyzed waveforms to make best use of the Synclavier (R) hardware in resynthesizing the timbre.

The default setting for USE RAW PHASES is "NO" (disabled).

HARMONIC NOISE FLOOR

The HARMONIC NOISE FLOOR parameter allows you to set a minimum threshold, below which harmonic information will be treated as noise and will be not used in the resynthesis of the frames. Occasionally, when trying to resynthesize inharmonic sounds, such as the piano, the Analysis Program will introduce a small amount of high frequency harmonic distortion. By setting the HARMONIC NOISE FLOOR to an appropriate value, such spurious harmonic coefficients can be suppressed. The range of values assignable to the HARMONIC NOISE FLOOR parameter goes from zero up to a maximum of 1000. These values correspond to harmonic coefficient strengths of 0.0 and 100.0 in the Synclavier (R) Real-Time system. Since the noise components would be quite small, typical SPECTRAL NOISE FLOOR settings are on the order of 1 or 2.

The default setting for HARMONIC NOISE FLOOR is 0.

Setting the Parameters for HORNF4

Now set the analysis parameters for the demonstration sound file. Since file HORNF4 has an OCTAVE setting of 4.05009, this same number should appear on the APPROXIMATE PITCH OF SOUND parameter line.. For now, do not change this value. Using the arrow keys, move the cursor to the PITCH WINDOW line and enter a value of 15%. Now move to the NUMBER OF HARMONICS parameter, and enter the number 28. The sample trumpet has a very wide bandwidth. The given setup will analyze out to approximately 20 kHz. For this first analysis, no pitch tracking will be used. And best results will be obtained by ignoring any phase variations in the analyzed frames. Leave the HARMONIC NOISE FLOOR parameter set at zero (0).

In summary, the analysis parameters should be set as shown below for the sound file HORNF4:

APPROXIMATE PITCH OF SOUND:	4.05009
PITCH WINDOW:	15%
USE RAW PITCH:	NO
NUMBER OF HARMONICS:	28
USE RAW PHASES:	NO
HARMONIC NOISE FLOOR:	0

PERFORMING THE ANALYSIS

Once the parameters for analysis are set, you press KP1 on the Analysis Main Menu to initiate the analysis, as stated above.

Throughout the analysis, you will receive feedback in the form of text and graphical displays which will guide you in evaluating the progress of any analysis.

Whenever the program locates a single Analysis label or a pair of Analysis labels, the screen will clear, and the name of the label or label pair, as well as their starting and ending times, will be written on the display. The current frame number will also be displayed. After the complete cycle has been extracted, either by reference to the two labels or by automatic cycle extraction, a plot of the raw waveform will be drawn on the screen. The Analysis Program will then proceed to compute the spectral properties of the extracted waveform. When this calculation has been completed, a plot of the waveform as it will look when resynthesized will be drawn. Next to these two plots will be drawn a plot of the harmonic magnitudes which were extracted from the sound file.

The Analysis Program has been designed so that you can start an analysis which can run to completion without requiring any further user intervention. The program will pause for approximately two seconds between frame analyses to give you an opportunity to view the results of the analysis for that frame. However, if you should want to view the results for a longer period of time, simply press the NO SCROLL key on the terminal to temporarily halt the analysis. When you are ready to proceed to the analysis of subsequent frames, press the NO SCROLL key again.

An additional aid that has been put into the Analysis Program is an elapsed time clock. During the analysis, a small clock display will be drawn in the lower left corner of the screen, showing how much time has elapsed since beginning the current analysis. Since some of the program computations are somewhat time-consuming, this frequently updated feedback is your assurance that the program is still running properly. As long as the clock continues to be updated on a fairly regular basis, all is well.

At any time during the analysis, you may press the BREAK key to stop analyzing. All frames which have already been processed will be preserved in memory. In this way, you can analyze parts of a timbre without having to analyze the entire sound.

Once an analysis has completed, or after you press BREAK, the screen will once again be cleared, and the Main Menu will be redrawn. At this time you would probably press the PF2 key to transfer control to the Resynthesis Release of the Synclavier (R) Real-Time System. After you press PF2, there will be a delay during which the Analysis Program will perform the final task of setting up the Synclavier (R) timbre definition. (Alternatively, you could have pressed KP1 to save the Analyzed timbre before going to the Real-Time Performance system.)

Analyzing HORNF4

You should now press KP1 to initiate analysis of the demonstration file HORNF4. When the analysis is completed and the Main Menu is redrawn, press PF2. After a brief pause, the digital display window on the keyboard should light up, and your resynthesized Synclavier (R) timbre will be active on the keyboard. It is a good idea to get in the habit of saving the timbre at this time, so that you will not lose the results of your analysis in the event that you accidentally destroy the keyboard timbre.

ADDING THE FINAL POLISH AT THE KEYBOARD

The trumpet timbre you have resynthesized is not quite right yet. There is no bite on the front of the sound, and there is no vibrato. You can use any standard Synclavier (R) parameter to complete the timbre.

Dial in the following settings on your keyboard:

Parameter	Value	Units
FM Ratio	.862	
HARMONIC ENVELOPE ATTACK	25	milliseconds
HARMONIC ENVELOPE INITIAL DECAY	50	milliseconds
HARMONIC ENVELOPE PEAK	13	
VIBRATO RATE	5.20	Hertz
VIBRATO DEPTH	.08	semitones
VIBRATO ATTACK	350	milliseconds

As you can see, just a few simple modifications to the Synclavier (R) timbre after it arrives on the keyboard fresh from the Analysis Program are all that is needed to put the final polish on the resynthesized sound. The Analysis Program has accomplished 95 percent of the timbre development task automatically. With the additional five percent provided by keyboard timbre programming, you have created a rich, live timbre, unlike any synthesized sound you have ever heard before.

A PRACTICAL GUIDE TO RESYNTHESIS

The following step-by-step guide to analysis/resynthesis has been prepared by Brian Banks and Anthony Marinelli, two synthesists with several years' experience with the Sample-to-Disk system. Their approach, which has produced highly successful results, may help you with your own analysis/resynthesis efforts.

1. Recall the sampled sound that you wish to resynthesize.

2. Change the horizontal scale so that the entire sound file is displayed.

3. Label the major points of change (e.g., ATTACK, DECAY1, DECAY2, SUSTAIN, etc.)

4. Set the horizontal scale at the default scale by typing SET HOR .01 and return the cursor to the first sample in the file by typing DISPLAY 0.

Labeling the Attack

Labeling the attack, which is usually no longer than 100 milliseconds, can be the most difficult aspect of resynthesis. First determine if the waveform is periodic or aperiodic in the attack portion.

If the waveform is periodic:

1. Place an Analysis label A_1 on the second sample in the file.

2. Type CEN * + .005. This centers the sample which is five milliseconds after the labeled sample.

3. Find the same point in phase as A_1 after this point and add label B_1.

4. Continue labeling start points in this manner to cover the initial part of the attack.

5. Subsequent labeling of the attack portion should occur between 10 to 50 milliseconds apart. Look for radical waveform or amplitude changes.

If the waveform is aperiodic:

To resynthesize aperiodic waveforms, you must specify an end point as well as a start point for a cycle to be analyzed correcty. The length of the waveform may be estimated by looking at a more periodic waveform in the steady state portion of the sound file.

1. Place an Analysis label A_1 on the second sample in the file.

2. Place label A_2 at the next 0 crossing that is at least one cycle long.

3. As long as there are new looking waveforms, keep labeling start and end points at 0 crossings.

Labeling the Next Sections

1. Set the horizontal scale to equal the length of the next major section (e.g., the first decay). Type DISPLAY DECAY1 to display the beginning of that section.

2. Label the points of major change within this section (e.g., a, b, c,...)

3. Set the horizontal scale to .01 again and display each label so that you may refine the rough labels and change them to Analysis labels.

4. Repeat these three steps for the other major sections.

Looking at the Steady State Tone

1. Display a steady state portion of the sound file, usually in the middle of the file. Starting from the default SET parameters, change the following:

   a. Type SET FFT 9000 (The program will automatically truncate this to 8192.)

   b. Type SET LEN .2 (The program will automatically truncate this to .163840.)

2. Type SPE

3. Move the cursor to the fundamental at its greatest magnitude and note the Frequency and the octave.pitch class (listed at the top of the right-hand box).

4. Push PF1 to display the harmonic locations and see if they line up with the peaks in the spectral display. If they do, great. If they don't, be aware that resynthesis will conform the non-harmonic overtones to the nearest coefficients which may or may not fool the ear. Filtering out the harmonic overtones from the non-harmonic overtones and resynthesizing each file

separately may yield better results.

Running the Analysis

1. Type SYN

2. Fill in the blanks for the analysis parameters as follows:

   Approximate Pitch of Sound:

   Type octave.pitchclass from step 3 above.

   Pitch Window:

   Specify 5 to 10% for steady tones. For vibrato and pitch bends, specify 20 % or more.

   Use Raw Pitch:

   If you wish to retain original pitch deviations, as in vibrato or wild attacks, select YES. However, be aware that tuning problems may occur. If you do not want pitch change, select NO.

   Number of Harmonics:

   Use the equation:

   Sample Rate ÷ Frequency ÷ 2 = Number of harmonics

   For example,

   50,000 ÷ 440 ÷ 2 = 57

   Use Raw Phase:

   If you have been careful to label the points nearest to 0, you can use phase information for more life-like results. If the near 0 points were not carefully selected, you may have a disaster; in which case, select NO.

   Harmonic Noise Floor:

   Type 0 to 2 if samples were taken direct to digital or to Winchester. If samples come from noisier sources, this number should increase accordingly.

3. Now press KP1 and wait . . .

Good luck.

Brian Banks
Anthony Marinelli